

This application is submitted in the name of inventors Jerrold Hauck and Colin
Whitby-Strevens.

SPECIFICATION

METHOD AND APPARATUS FOR BORDER NODE BEHAVIOR ON A FULL-DUPLEX BUS

BACKGROUND OF THE INVENTION

[0001] This application is a divisional of U.S. patent application serial number 09/484,479, filed 1/18/2000, incorporated by reference herein in its entirety.

Field of the Invention

[0002] The present invention relates to bus management. In particular, the present invention relates to the behavior of border nodes within a high performance serial bus system.

The Prior Art

Background

[0003] Modern electronic equipment has greatly enhanced the quality of our lives. However, as the use of such equipment has increased, so has the need to connect equipment purchased from different manufacturers. For example, while a computer and a digital camera may each be useful when used alone, the ability to connect the digital camera to the computer and exchange information between the two makes the combination even more useful. Therefore, a need was apparent for a serial bus standard that would allow for the connection and communication between such devices.

[0004] The IEEE 1394-1995 standard was developed to satisfy this need. This standard revolutionized the consumer electronics industry by providing a serial bus management system that featured high speeds and the ability to “hot” connect equipment to the bus; that is, the ability to connect equipment without first turning off the existing connected equipment. Since its adoption, the IEEE 1394-1995 standard has begun to see acceptance in the marketplace with many major electronics and computer manufacturers providing IEEE 1394-1995 connections on equipment that they sell.

[0005] However, as technologies improved, the need to update the IEEE 1394-1995 standard became apparent. Two new standards are being proposed at the time of the filing of this application, herein referred to as the proposed IEEE 1394a, or P1394a standard, and the proposed IEEE 1394b, or P1394b standard. Improvements such as higher speeds and longer connection paths will be provided.

[0006] In the discussion that follows, it will be necessary to distinguish between the various standards that are being proposed as of the date of this application.

Additionally, it will be necessary to distinguish hardware and packet transmissions that are compatible with the P1394b standard and not earlier standards.

[0007] Thus, the term “Legacy” will be used herein to refer to the IEEE 1394-1995 standard and all supplements thereof prior to the P1394b standard. Thus, for example, a Legacy node refers to a node compatible with the IEEE 1394-1995 standard and all supplements thereof up to, but not including, the P1394b standard.

[0008] Additionally, packets of data will be referred to herein depending on the context the packets are in. For example, a packet of data that is compatible with the P1394b standard and is travelling through a PHY compatible with the P1394b standard will be referred to as Beta format packets. Packets of data that are compatible with the Legacy standard but are travelling through a PHY compatible with the P1394b standard will be referred to as Legacy format packets. Finally, packets of data that are compatible with the Legacy format and are travelling across a data strobe link will be referred to as Alpha format packets.

[0009] Furthermore, in the discussion that follows PHYs that are compatible with the P1394b standard may be referred to in various ways, depending upon the context the PHY is operating in and the capability of the PHY. For example, a PHY that has circuitry compatible with the P1394b standard but not any previous standards will be referred to as a B only PHY. Also, a PHY that is compatible with the P1394b standard and is directly attached with only devices compatible with the P1394b standard will be referred to as B PHYs. Finally, a PHY that is communicating with both Legacy devices

and devices compatible with the P1394b standard will be referred to as a border device, border PHY, or border node.

[0010] Finally, a communications systems that has only B PHYs attached will be referred to as a B bus.

Data transmission in Legacy systems

[0011] One area that has been improved in the P1394b standard is in the way that data transmission takes place on the bus.

[0012] Figure 1 is a prior art example of a Alpha format data packet 100 according to Legacy specifications. In the Legacy specifications, a data packet will begin with the transmission of a Data Prefix (“DP”) identifier, shown as DP 102 in Fig. 1. Importantly, in the Legacy specification , a DP must have a duration of no less than 140 nanoseconds (ns), though a DP may be of any greater length.

[0013] Typically, a DP is followed by the transmission of clocked data, known as the payload, shown as clocked data 104 in Fig. 1. On a Legacy bus, the payload will be clocked at a rate of 100 Megabits per second (Mb/s), 200 Mb/s, or 400 Mb/s. These data rates are known as S100, S200, and S400, respectively.

[0014] Finally, the payload is followed by a Data End (“DE”), shown as DE 106 in Fig. 1. In the Legacy specifications, a DE must be at least 240 ns in length.

[0015] As is appreciated by one of ordinary skill in the art, the Legacy specifications thus define a timer-based system, where data transmission begins and ends according to a fixed timer.

Compatibility issues in Legacy systems

[0016] As mentioned above, there are three clocked data rates present in Legacy systems, S100, S200, and S400. Initially, when the IEEE 1394-1995 standard was introduced, devices could only communicate at the S100 rate. Later, devices were introduced that communicated at the S200 and S400 rates.

[0017] One problem that occurred in the prior art was how to insure compatibility between the various devices on the market that were communicating at these different rates.

[0018] Figure 2 illustrates such a compatibility problem. Fig. 2 has three nodes, nodes #0, #1, and #2. Node #2, the root node in this example, wishes to communicate with node #1. As is indicated in Fig. 2, nodes #1 and #2 are capable of communicating at the S400 data rate, while node #0 is only capable of communication at the lower S100 rate.

[0019] Figure 3 illustrates the prior art solution of speed filtering on a Legacy bus. Fig. 3 shows root node #2 transmitting S400 data in packet P1 to node #1. In the prior art, to prevent node #0 from receiving the S400 data that it cannot understand, it is “shielded” from such data by having root node #2 transmit a null packet P2 to it.

[0020] In the Fig. 3 Packet Detail illustration, packets P1 and P2 are shown together on a common time axis. Packet P1 comprises a DP 300, S400 data 302, and a DE 304. Null packet P2 comprises a DP 306, and a DE 308. As is appreciated by one of ordinary skill in the art, the null packet accomplishes its shielding by extending the DP for the amount of time required to send S400 data 302. As is known by those of ordinary skill in the art, on a Legacy bus all nodes must remain synchronized in their interpretation of idle time. Thus, the null packet effectively 'busies' node #0 while root node #2 transmits S400 data to node #1 and thus shields node #0 from speeds it cannot understand.

Data transmission in P1394b

[0021] Figure 4 is a representation of the prior art data packet structure according to the P1394b standard. As is known by those of ordinary skill in the art, P1394b utilizes a packet structure that is scaled to speed unlike the fixed timer system utilized in Legacy standards. Specifically, the packet structure in P1394b is based upon symbols, rather than the fixed intervals found in the Legacy standards.

[0022] In Fig. 4, a typical prior art packet 400 of P1394b data is shown. As is known by those of ordinary skill in the art, a data packet begins in P1394b with the transmission of at least two packet starting symbols. In this example, a Speed Code symbol Sb1 and a Data Prefix symbol DP1 are shown as the packet starting symbols. Then, P1394b data bytes B1 through Bn are transmitted. Bytes B1 through Bn may be

referred to herein as the payload. Finally, the transmission of data is terminated by transmitting DE symbols DE1 and DE2.

Compatibility problems between P1394b and Legacy nodes and clouds

[0023] Figure 5 shows a data communications system comprising both P1394b and Legacy devices. This type of system will be referred to herein as a “hybrid” system. Fig. 5 shows a Legacy node #0 connected to Legacy node #4, forming what is known as an “Legacy cloud” of Legacy nodes.

[0024] Legacy node #4 is connected to border node #3. Border node #3 is connected to B PHYs #1 and #2, forming what is known as a “Beta cloud” of border and B PHYs.

[0025] One problem encountered with systems containing both Legacy and P1394b compliant nodes is how to ensure that the newer P1394b PHYs know that there are older Legacy devices present on the bus. Hence there is a need for a method to determine the existence of a hybrid bus such as that of Fig. 5.

[0026] Furthermore, given a hybrid bus, there is a need to establish one border node as the senior border node within a given cloud. Furthermore, given a hybrid system with many clouds, there is a need to establish one border node from each cloud as a senior border node.

[0027] Another problem raised by the configuration of Figure 5 is that of arbitration on a hybrid bus. As is known by those of ordinary skill in the art, the Legacy

standard employs a single request type for both asynchronous and isochronous arbitration. The first request to be received by the root is granted immediately, and all other requests are denied and required to be withdrawn. P1394b, however, uses a pipelined arbitration system, pipelining requests for both the asynchronous and isochronous bus phases. Thus the BOSS node, which is responsible for granting arbitration requests in P1394b, nominally must be aware of the phase of the bus in order to properly grant arbitration requests.

[0028] However, there may be a case where the BOSS node is unaware of when the bus enters into an isochronous phase. For example, this may happen when there are no isochronous-aware nodes within the Beta cloud that the BOSS is in, while a Legacy cloud is introducing an isochronous request into the Beta cloud. In these cases, the nodes in the Beta cloud will continue to assume that the phase of the bus is asynchronous, and the grant for the Legacy isochronous request may be handled incorrectly. Therefore, there is a need for an arbitration protocol for properly handling arbitration requests throughout multiple clouds.

[0029] Another problem raised by the configuration shown in Fig. 5 is that of managing the timers required by the Legacy standard. As mentioned above, the Legacy standard requires that bus protocols occur according to a gap timer-based system, whereas the P1394b standard abandons the timers in favor of a symbol-based system.

[0030] Border devices already contain the necessary logic to manage gap timers, since by definition border nodes have Legacy compatible ports. However, B only

devices do not contain the circuitry necessary to manage timers according to the Legacy standard. Therefore, there is a need for a protocol for freeing B only devices of the gap timer requirements and transferring this responsibility to border nodes, thus freeing B only devices from the requirement of having gap timer logic built in.

[0031] Another issue raised by the hybrid system shown in Figure 5 is what happens if the bus fails. As is known by those of ordinary skill in the art, occasionally an acknowledge or a grant can be lost in the system, throwing the system into a “race” state, whereby no device is in charge of the system. In other words, a condition can occur where no device knows the current state of the system. Currently, there exists no protocol for a hybrid bus to prevent or cure race conditions. Hence, there is a need for a protocol to allow a hybrid bus to return to a known state.

[0032] A final issue concerning the system of Fig. 5 is related to what is known as “quiet times” in the Legacy standard. As is known by those of ordinary skill in the art, in order to arbitrate for control of the bus, a Legacy node must sometimes wait for a subaction or arbitration reset gap to occur. However, in order to ensure that all nodes on the bus see the same subaction or arbitration reset gap, hysteresis is built into the process by making a node wait an additional period known as an “arb_delay” before arbitrating. This arb_delay guarantees that in a worst case round-trip scenario across the bus, every node has seen the subaction gap or the arbitration gap. The period comprising the subaction gap plus the arb_delay is known as quiet time #1, and the period comprising the arbitration reset gap plus the arb_delay is known as quiet time #2, and no arbitration requests may be made during either quiet time.

[0033] There is a danger in a hybrid bus of a B PHY granting requests during this quiet time. Currently, there is no protocol for insuring that B PHYs do not interfere with Legacy arbitration timing schemes, other than forcing B PHYs to adhere to Legacy standards, which is undesirable since the performance gains inherent in the P1394b standard are lost. Hence, there is a need for a protocol which allows the coexistence of Legacy and P1394b arbitration in a hybrid bus.

BRIEF DESCRIPTION OF THE INVENTION

[0034] The invention satisfies the above needs. The present invention relates to bus management. In particular, the present invention relates to the behavior of border nodes within a high performance serial bus system.

[0035] A method for determining and communicating the existence of a hybrid bus is disclosed, comprising the acts of: determining whether the node is a border node; forwarding isochronous and asynchronous requests as normal if the node is not a border node; and if the node is a border node, then; determining whether the node has any asynchronous requests to forward; issuing a Border_low request if there are no asynchronous requests to forward; forwarding the asynchronous requests if there are asynchronous requests to forward; determining whether there are any isochronous requests to forward; issuing a Border_low request if there are no isochronous requests to forward; and forwarding the isochronous requests if there are isochronous requests to forward.

[0036] A method for determining and communicating the existence of a hybrid bus is disclosed, comprising the acts of: determining whether the node has a connection to a Legacy link layer; if the node determines that it has a connection to a Legacy link layer, then transmitting a Self-ID packet without a Speed Code; and if the node determines that it does not have a connection to a Legacy link layer, then transmitting a Self-ID packet with a Speed Code.

[0037] A method for determining and communicating the existence of a hybrid bus comprising the acts of: examining received Self-ID packets by the P1394b compliant node for the absence of a Speed Code; and presuming the existence of a hybrid bus if any of the received Self-ID packets do not contain a Speed Code.

[0038] A method for determining a path to a senior border during the Self-ID process is disclosed, comprising the acts of: marking the border node as the senior border; determining whether the border node has received a Self-ID packet that does not contain a Speed Code on a parent beta port of the border node; marking the port on the border node as the path to the senior border node if the border node has received a Self-ID packet that does not contain a Speed Code on a parent beta port of the border node; and canceling the border node's own status as a senior border node.

[0039] A method determining a path to a senior border is disclosed, comprising the acts of: determining whether a B PHY has received a Self-ID packet without a Speed Code on at least one port; and marking the last at least one port on the B PHY to receive a Self-ID packet without a Speed Code as the path to the senior border and canceling by the B PHY of any other ports within the node that have been marked as the path to

the senior border node if the B PHY has received a Self-ID packet without a Speed Code on at least one port.

[0040] A method for identifying a senior border node is disclosed, comprising the acts of: determining whether the border node was the last node within the cloud to transmit a Self-ID packet; and marking the node as the senior border node if the node was the last node within the cloud to transmit a Self-ID packet.

[0041] A method for returning control to the senior border node is disclosed, comprising: granting BOSSship to a predetermined port by the beta device; and in the alternative, passing control of the system towards the senior border node by the beta device.

BRIEF DESCRIPTION OF THE DRAWING FIGURES

[0042] Figure 1 is a prior art Alpha Format Packet example.

[0043] Figure 2 is a prior art Legacy Compatibility example.

[0044] Figure 3 is a prior art Legacy Speed Filtering example.

[0045] Figure 4 is a prior art P1394b Packet format example.

[0046] Figure 5 is a prior art Data Communications System having both P1394b and Legacy devices.

[0047] Figure 6 is a Flowchart For Detection And Communication Of Hybrid bus According To The Present Invention.

[0048] Figure 7 is a prior art Legacy Format Packet.

[0049] Figure 8 is a Flowchart For Detection And Communication Of A Legacy Link Layer According To The Present Invention.

[0050] Figure 9 is a Flowchart For Detection And Communication Of A Hybrid Bus According To The Present Invention.

[0051] Figure 10 is a flowchart for Detection of Path to Senior Border By a Border PHY According to the Present Invention.

[0052] Figure 11 is a flowchart for Identification of Senior Border Node By a B PHY According to the Present Invention.

[0053] Figure 12 is a flowchart for Restoring Senior Border Node as BOSS According to the Present Invention.

[0054] Figure 13 is a flowchart for an OK_TO_GRANT Profile According to the Present Invention.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

[0055] Persons of ordinary skill in the art will realize that the following description of the present invention is illustrative only and not in any way limiting. Other

embodiments of the invention will readily suggest themselves to such skilled persons having the benefit of this disclosure.

[0056] The present invention relates to data communications. More particularly, the present invention relates to a method and apparatus for an arbitration and fairness protocol on a serial bus. The invention further relates to machine readable media on which are stored embodiments of the present invention. It is contemplated that any media suitable for retrieving instructions is within the scope of the present invention. By way of example, such media may take the form of magnetic, optical, or semiconductor media.

[0057] The present invention relates to data structures and the transmission of such data structures. It is contemplated that the present invention may be embodied in various computer and machine readable data structure. Furthermore, it is contemplated that data structures embodying the present invention will be transmitted across computer and machine readable media.

[0058] The present invention may be described through the use of flowcharts. Often, a single instance of an embodiment of the present invention will be shown. As is appreciated by those of ordinary skill in the art, however, the protocols and procedures described herein may be repeated continuously or as often as necessary to satisfy the needs described herein. Accordingly, the representation of the present invention through the use of flowcharts should not be used to limit the scope of the present invention.

[0059] The present invention further relates to devices that embody the P1394b standard. By way of example, such devices may include those typically used in an audio/video entertainment system, such as home theater receivers, DVD players, computers, or hand-held devices such as cameras and the like. The devices may also include those industrial in nature, such as test and measurement equipment, professional audio/video recording devices, as well as system control or robotic devices found in an industrial environment.

[0060] The invention also relates to nodes and physical computers, such as state machines. The present invention may be embodied in any collection of nodes linked together through a bus. Typically, each device connected to the bus will also have one corresponding node physical layer controller embedded therein. However, a given device may have more than one node, and therefore it follows that one device may have more than one connection to more than one bus. For the discussion that follows, the examples will show the typical situation where one node corresponds to one device.

[0061] Each node may communicate to other nodes in an P1394b-compatible system through links. Typically, a cable is used for a link, as is provided for in the P1394b standard. However, any communication means may be employed. By way of example, an infrared, RF, or other wireless system may be used, as well as an optical system.

[0062] Typically, a link is coupled to a node through a port. A port transmits and receives messages and data between the node and link. As is known by those of ordinary skill in the art, each node may have more than one port.

Detecting the presence of a hybrid bus

[0063] As mentioned in the prior art discussion above, one problem associated with the prior art was the need for a method of identifying a hybrid bus. Two methods according to the present invention for detecting a hybrid bus will now be disclosed.

Detection through the use of P1394b request symbols

[0064] A first preferred embodiment takes advantage of certain characteristics inherent in the P1394b standard to detect the presence of a Legacy device or a Legacy link layer. As is known by those of ordinary skill in the art, when a P1394b PHY device is connected to a bus, the P1394b PHY will cycle through certain start-up procedures, known as the “up-start” procedure. One of the by-products of the up-start procedure is that a P1394b PHY will have knowledge of all of the connections that have been made to its ports when the up-start procedure is finished. As a result, at the conclusion of the up-start procedure, if a node has both Legacy and P1394b connections, it will know that is a border node.

[0065] As is known by those of ordinary skill in the art, arbitration in the P1394b standard occurs through the use of requests that are pipelined to the BOSS node according to priorities. Furthermore, the protocols in P1394b provide that a node will forward the highest priority request it hears on to the BOSS node. In order to communicate knowledge of the existence of a hybrid bus to the rest of the bus, a special symbol known as Border_low was developed.

[0066] Figure 6 shows the detection and communication of a hybrid bus according to the present invention. In query 600, a node first determines whether or not it is a border node. This detection may occur in any manner, such as the up-start procedure described above.

[0067] If the node is not a border node, then in act 601 the node will forward any asynchronous and isochronous requests present as normal, and the process ends.

[0068] If the node is a border node, then in query 602 the node next determines whether there are any asynchronous requests to forward on this port. If there are asynchronous requests pending, the node will drain those requests as shown in act 603. If there are no asynchronous requests to forward, then the node will issue a Border_low request as shown in act 604.

[0069] Next, in query 606, the node determines whether there are any isochronous requests to forward on this port. If there are isochronous requests pending, the node will forward those requests as shown in act 605. If there are no isochronous requests to forward, then the node will issue a Border_low request as shown in act 608.

[0070] In a preferred embodiment of the present invention, the Border_low request is given a low enough priority to allow normal arbitration to take place first. Then, when there are no other in-phase arbitration requests pipelined, all other nodes will forward the Border_low request towards the BOSS node, thus communicating to the BOSS node and those nodes along the path towards the BOSS node that there is a Legacy device on the bus.

[0071] In another preferred embodiment of the present invention, there are separate Border_low request symbols for both the asynchronous and isochronous arbitration phases.

[0072] In one embodiment, the Border_low request has been encoded within data packets transmitted on a P1394b-compliant bus. In another embodiment, the Border_low request has been programmed and encoded in logic on a PHY.

Detection of a hybrid bus through the use of the Self-ID sequence

[0073] As is known by those of ordinary skill in the art, when a node initializes there is a Self-ID stage wherein each node transmits its identity and certain characteristics about itself to the rest of the bus. As is further known by those of ordinary skill in the art, when a node transmits a packet on a bus, there is a speed code contained therein that identifies the clock rate of the data being transmitted. The S100 Alpha format packet contains no speed code, and thus no speed code is repeated into the Beta cloud. The present invention utilizes both of these characteristics.

[0074] Figure 7 shows an example of the beginning of a prior art Legacy packet format. Within packet 700, there is first a data prefix DP1 and a Speed Code symbol. As is known by those of ordinary skill in the art, the Speed Code symbol will contain an indication therein of the clock rate of the data that will follow DP2 in packet 700.

[0075] Figure 8 is a flowchart of detection and communication of a Legacy link layer according to the present invention. In query 800, the PHY will determine if has a

Legacy link layer attached to it. If the node determines that it has a Legacy link layer attached, then it will transmit a Self-ID on the bus that has no Speed Code and presume a hybrid system as shown in act 802. If the node determines that it does not have a Legacy link layer attached, then it will transmit a Self-ID on the bus with a Speed Code in act 804.

[0076] Figure 9 is a flowchart of the detection and communication of a hybrid bus according to the present invention. In query 900, all P1394b PHYs will examine each Self-ID that it receives, and will look for the absence of a Speed Code. If there is no Speed Code in a received Self-ID packet, then in act 902, the node will presume that the Self-ID packet passed through a Legacy link some where in the chain, or that a Legacy link layer is present, thus indicating a hybrid bus. In a preferred embodiment, the node will store this presumption until the next bus reset.

[0077] In one embodiment of the present invention, this storing is accomplished by way of setting a designated bit within logic in the PHY. In another embodiment of the present invention, this storing is accomplished through the use of a variable in which the state of the bus is stored.

Establishing a senior border node

[0078] As mentioned in the prior art section above, there is a need to establish a senior border node among border nodes in one or more clouds. The following discussion discloses a method for determining the path to, and the identification of, a senior border node.

[0079] As is known by those of ordinary skill in the art, nodes assign themselves a number, known as a physical ID or a PHY-ID, during the Self-ID process with the lowest numbers being assigned first. It is contemplated that the process according to the present invention may take place at any time after a bus reset and before the completion of the Self-ID process.

[0080] Figure 10 is a flowchart of the detection of the path to the senior border node by a border PHY according to the present invention.

[0081] As is known by those of ordinary skill in the art, in the Self-ID process a node will first receive Self-ID packets from its children, if any. The node will then generate its own Self-ID packet. Then, the node will receive Self-ID packets from its parent port.

[0082] The following process describes how a particular border node will detect the path to the senior border.

[0083] In act 1000, the border node will mark itself as the senior border node. It is contemplated that this act takes place any time after bus reset. In query 1002, the border node will then determine whether it has received a packet on its parent beta port that does not contain a Speed Code during the Self-ID process. It is contemplated that the detection may take place through different methods. By way of example, in one preferred embodiment of the present invention, all of the received Self-ID packets are stored in a queue and the detection for the absence of a Speed Code is performed on all received packets at once. In another preferred embodiment of the present invention, the

detection for the absence of a Speed Code takes place at the time each Self-ID packet is received.

[0084] If the border node has received a packet on its parent beta port that does not contain a Speed Code during the Self-ID process, then the node marks its parent port as path towards senior node and cancels its own status as senior border node in act 1004. If the border node fails to receive a packet on its parent beta port that does not contain a Speed Code during the Self-ID process, then the process ends.

[0085] In one embodiment of the present invention, this marking is accomplished by way of setting a designated bit within logic in the PHY. In another embodiment of the present invention, this marking is accomplished through the use of a variable in which the location of, or the path to, the senior border node is stored.

[0086] Figure 11 shows the identification of the path to the senior border node by a B PHY according to the present invention. In query 1100, the B PHY determines whether it has received a Self-ID packet without a Speed Code. If it has received a Self-ID packet without a Speed Code, then the B PHY marks the last port to receive Self-ID packet without a Speed Code as the path towards senior border and cancels any other ports having this status. It is contemplated that this test will take place during the Self-ID process, and may be repeated as often as necessary to satisfy the needs described herein. It is further contemplated that this marking may take place through the means described above.

[0087] If the B PHY has not received a Legacy Self-ID packet in query 1100, then the process ends.

Properly carrying out arbitration requests across multiple clouds

[0088] As mentioned in the prior art section above, there is a need for an arbitration protocol for properly handling arbitration requests throughout multiple clouds. The following discussion will disclose a protocol whereby a border node may properly forward an arbitration request received from an Legacy cloud.

[0089] As is known by those of ordinary skill in the art, when an Legacy device requests arbitration, it may only do so at a time when it is proper. This is due to the strict arbitration protocol of the Legacy standard. Therefore, if a border node receives an arbitration request from a Legacy device or link, the border node may safely assume that the request was made properly in relation to the quiet times describe above.

[0090] However, as mentioned in the prior art section above, it is possible that a pipelined P1394b request may not be properly granted in relation to the quiet times within the Beta cloud because of the dual-phase approach used in P1394b. To solve this problem, the present invention introduces a new request, known as a Legacy request. In a preferred embodiment of the present invention, the Legacy request is deemed to be a current request and will be given a higher priority than both asynchronous and isochronous requests. Thus, the Legacy request will be forwarded to the BOSS node unimpeded and immediately granted.

[0091] According to a preferred embodiment of the present invention, when a border node receives an arbitration request from a Legacy cloud, the border node will “convert” the request into a Legacy request and forward the Legacy request to the current BOSS. Thus, because of this new type of request, a B PHY which is unaware of the current bus phase may correctly service a Legacy request without interrupting the current bus phase by incorrectly granting an out-of-phase request.

[0092] In one embodiment, the Legacy request has been encoded as a request symbol transmitted on a P1394b-compliant bus. In another embodiment, the Legacy request has been programmed and encoded in logic on a PHY.

Freeing B only PHYs from Legacy gap timers

[0093] As mentioned in the prior art discussion, there is a need for a protocol that transfers the gap timing requirement away from B only devices and to border devices, which already have the necessary logic built into them. The following discussion will now disclose such a protocol.

[0094] In a preferred embodiment of the present invention, when one or more border devices is present in a system, B PHYs will refrain from issuing gap tokens, and the border devices will take over the responsibility of timing the IDLE duration. Gap tokens are the BOSS equivalent to a Legacy gap event. This means that when a border device detects a duration which amounts to a subaction_gap, the border will issue an Asynch_start. Likewise, when the border device detects a duration amounting to an Arb_reset_gap, the border device will issue an Arb_reset_even/odd.

[0095] The above protocol may be implemented in a number of ways. For example, in one embodiment, all border devices automatically time the gaps and issue an event gap token when their individual timers expire.

[0096] In another embodiment, the border devices will utilize a low-priority arbitration to select one of the border devices to ensure that one border device is selected as BOSS before an extended period of IDLE appears on the bus.

[0097] In a preferred embodiment, the senior border as disclosed above, is given responsibility for issuing gap tokens in the Beta cloud. Normally, when a P1394b PHY has finished all of its tasks, it would issue a phase change. However, in this embodiment, when a P1394b PHY is finished with its tasks, it turns over control to the senior border node. The senior border can then ensure compliance with the gap timers. Recall from the discussion of the senior borders that all P1394b PHYs within the Beta cloud will know the path to their senior border node.

Restoring the local root as BOSS

[0098] In a preferred embodiment of the present invention, when a hybrid bus is present, control within a Beta cloud will be passed to the senior border node within the cloud. Thus by passing control to the senior border node, the Beta cloud is assured that arbitration will be performed correctly, even if the senior border node has a Legacy cloud as a parent.

[0099] To accomplish this, the present invention discloses a protocol in which a BOSS device which is transmitting immediately communicates one of two intentions regarding its BOSS status, or “BOSSship”, at the end of its transmission: 1) either explicitly grant BOSSship to a specific port, or 2) pass control towards the senior border node. By passing control towards the senior border node, this protocol insures that eventually control will revert to a device which may proxy between clouds in a hybrid bus. Thus, after transmitting, there is no longer a question as to who has control of the bus, and race conditions are thus prevented.

[0100] A preferred embodiment of the present invention is implemented through the use of new packet-ending symbols, shown in Table 1. In one embodiment, these symbols have been encoded and transmitted on a bus. In another embodiment, these symbols have been programmed and encoded in logic on a PHY.

Table 1 – Packet-ending symbols according to the present invention

<u>Symbol Name</u>	<u>Comment</u>
GRANT	The end of a subaction has been reached, and the receiving node is being elected BOSS.
DATA_NULL	The bus is being actively granted to another PHY, and a packet will follow. Therefore, the receiving node is not being elected BOSS and must not let its IDLE timers time.
DATA_END	When sent out a senior port: The end of a subaction has not been determined, and the receiving node connected to the senior port is being elected BOSS; and When sent out a junior port: the receiving node connected to the junior port is not being elected BOSS,

	should start or continue timing the IDLE gap.
--	---

[0101] Figure 12 shows how the symbols in Table 1 are utilized. In query 1200, the current BOSS determines if an end of subaction (“EOS”) has been reached. If an EOS has been reached, then in query 1202, the current BOSS determines if there are any in-phase requests to grant. If there are in-phase requests to grant, then in act 1204 the current BOSS sends a GRANT symbol to the granted port, and a DATA_NULL symbol to all other ports. If there are no in-phase requests to grant, then in query 1203 the node determines whether it is the senior border node. If the node is the senior border node, then in act 1207 the node send a Data End out all ports.

[0102] If the node is not the senior border node, then in act 1206 the current BOSS sends a GRANT symbol out its senior port, and a DATA_END symbol out its junior ports. Thus, this process achieves the goal of returning BOSSship towards the senior border node, while communicating to the receiving node connected to the senior port that an EOS was achieved.

[0103] Referring still to Fig. 12, if no EOS was reached in query 1200, then in act 1208 the current BOSS sends a DATA_END symbol out all ports. This returns control to the senior node, while allowing the IDLE timers to continue.

[0104] Thus, in an example where an acknowledge packet is never received, control will be returned to the local senior border node, which has the capability to

perform any necessary recovery. This recovery can occur even if the senior border node has a Legacy node as a parent.

BOSS OK to Grant

[0105] As mentioned in the prior art section, there is a need for a protocol to ensure compatibility when arbitration takes place on a hybrid bus. Such a protocol will now be discussed.

[0106] A protocol according to the present invention comprises two variables: BOSS and OK_TO_GRANT. Prior to issuing a grant, the BOSS will check the status of each of these variables. In a first embodiment of the present invention, only if both are true will the BOSS be able to issue a grant for pipelined P1394b request. In a second embodiment of the present invention, the BOSS may grant a Legacy request when only the BOSS variable is set.

[0107] In a preferred embodiment, these variables have been programmed and encoded in logic on a PHY.

BOSS

[0108] In the present invention, the BOSS variable indicates whether the BOSS has BOSSship – the right to issue a grant. The BOSS indicator can be moved freely within a cloud between nodes. The BOSS indicator can be set in one of two ways: 1) a P1394b-compliant node which receives a GRANT either implicitly or explicitly; or 2) a node which first transmits a packet into a Beta cloud automatically becomes BOSS of

that cloud. An explicit grant occurs by way of receiving a grant symbol, while an implicit grant occurs by way of receiving a DATA_END symbol from a junior port.

[0109] Furthermore, the present invention relates to a protocol which determines when a BOSS device must surrender its BOSSship by resetting its BOSS indicator. In one embodiment, whenever a packet is received from a P1394b port, the receiving PHY ceases to be BOSS. In another embodiment, whenever an implicit or explicit grant is issued, the issuing PHY ceases to be BOSS.

[0110] Thus, the BOSS indicator indicates whether the BOSS has acquired the exclusive right to issue a grant, either explicitly or implicitly. However, just because the BOSS has the right to issue a grant does not necessarily mean that it is safe to do so within a hybrid bus.

OK TO GRANT

[0111]Figure 13 is a profile of an OK_TO_GRANT indicator according to the present invention. Fig. 13 has three rows indicating the status of the bus at a certain moment, labeled End of Isoch (end of the isochronous packet), !EOS (not end of asynchronous subaction), and EOS (end of asynchronous subaction). Fig. 13 further has two columns labeled B bus, and hybrid bus.

[0112] At the intersection of the columns and rows in Fig. 13 are waveform representations of the OK_TO_GRANT indicator. When the waveform is represented as

being high, the OK_TO_GRANT indicator is set, and when the waveform is represented as being low, the OK_TO_GRANT indicator is reset.

B Bus/End of isoch

[0113] Referring still to Fig. 13, a case representing a B bus after the end of an isochronous packet is shown according to the present invention. On a B bus, and at the end of an isochronous packet represented as the EOP time marker on the waveform, there is no need to wait for an acknowledgment. As soon as the EOP arrives, the OK_TO_GRANT indicator will be set. Therefore, if a node has its BOSS indicator set, at the end of an isochronous packet the node may issue a grant according to the present invention.

B Bus/!EOS

[0114] This example represents the case where an EOP has arrived that was not observed to be the end of a subaction. By way of example, this may occur when the end of a primary packet has been marked with a data end. This illustrates the example where an ACK is expected but is missing. On a B bus, the BOSS must wait some predetermined amount of time before it may issue a grant, just in case an ACK arrives. This time period is represented by the time period shown between the EOP and ACK missing time markers prior to the OK_TO_GRANT being set. According to a preferred embodiment of the present invention, after the predetermined amount of time, the BOSS will assume that the ACK is missing, and will set its OK_TO_GRANT indicator.

B Bus/EOS

[0115] Referring still to Fig. 13, a case representing a B bus after the end of a subaction is shown according to the present invention. By way of example, this may occur where the end of a packet has been marked with a GRANT. In this case, at the end of a packet, the OK_TO_GRANT will be set according to a preferred embodiment of the present invention.

Hybrid Bus/End of isoch

[0116] Referring now to Fig. 13 in the Hybrid Bus column, a case representing a hybrid bus after the end of an isochronous packet is shown according to the present invention. After the end of an isochronous packet in a hybrid bus, there is an instant in time where the BOSS may safely issue a grant. This instant in time is represented in Fig. 13 by the momentary setting of the OK_TO_GRANT indicator at the EOP time marker. After this instant, the OK_TO_GRANT indicator is reset until the end of the subaction gap plus the arb_delay (quiet time #1). This ensures protection of the quiet time #1. Then, after quiet time #1, the BOSS momentarily may safely issue a grant. This is represented by a momentary setting of the OK_TO_GRANT indicator at the time marker labeled SAG. During the time period immediately after the momentary setting of the OK_TO_GRANT indicator, after the subaction gap and the arb_delay, and immediately up to, but not including the arbitration reset gap and the arb_delay (quiet time #2), the BOSS may not safely issue a grant. This is shown by the OK_TO_GRANT indicator being reset during quiet time #2. Finally, after quiet time #2, the BOSS may safely issue a

grant. This is represented by the OK_TO_GRANT indicator being set at the time marker labeled ARG.

Hybrid Bus/!EOS

[0117] Referring still to Fig. 13, a case representing a hybrid bus where the end of an subaction is missing is shown according to the present invention. When the end of an subaction is missing in a hybrid bus, the BOSS may not safely issue a grant, and must wait through the end. Then, after quiet time #1, the BOSS momentarily may safely issue a grant. This is represented by a momentary setting of the OK_TO_GRANT indicator at the time marker labeled SAG. During the time period immediately after the momentary setting of the OK_TO_GRANT indicator, after the subaction gap and the arb_delay, and immediately up to, but not including the arbitration reset gap and the arb_delay (quiet time #2), the BOSS may not safely issue a grant. This is shown by the OK_TO_GRANT indicator being reset during quiet time #2. Finally, after quiet time #2, the BOSS may safely issue a grant. This is represented by the OK_TO_GRANT indicator being set at the time marker labeled ARG.

Hybrid Bus/EOS

[0118] Referring still to Fig. 13, a case representing a hybrid bus after the end of an subaction is shown according to the present invention. After the end of a subaction in a hybrid bus, the BOSS may safely issue a grant during quiet time #1. This is one exception to the general rule of not violating the quiet times. This is similar to the accelerated arbitration in the P1394a standard. This is indicated by the setting of the

OK_TO_GRANT indicator at time marker EOP through time marker SAG. During the time period immediately after the setting of the OK_TO_GRANT indicator, after the subaction gap and the arb_delay, and immediately up to, but not including the arbitration reset gap and the arb_delay (quiet time #2), the BOSS may not safely issue a grant. This is shown by the OK_TO_GRANT indicator being reset during quiet time #2. Finally, after quiet time #2, the BOSS may safely issue a grant. This is represented by the OK_TO_GRANT indicator being set at the time marker labeled ARG.

[0119] While embodiments and applications of this invention have been shown and described, it would be apparent to those skilled in the art that many more modifications than mentioned above are possible without departing from the inventive concepts herein. The invention, therefore, is not to be restricted except in the spirit of the appended claims.